

APPLICATION FOR UNITED STATES LETTERS PATENT

METHOD AND APPARATUS FOR EFFICIENT LINK REDUNDANCY

Inventor:

Jim Mao

Prepared by:

Blakely, Sokoloff, Taylor & Zafman LLP
12400 Wilshire Boulevard, Seventh Floor
Los Angeles, California 90025-1026
(408) 720-8598

Docket No. 005043.P008

"Express Mail" mailing label number EL672751513US
Date of Deposit 12/22/00

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

MARA E. BROWN
(Typed or printed name of person mailing paper or fee)

[Signature]
(Signature of person mailing paper or fee)

09745689 "122200

Method and Apparatus For Efficient Link Redundancy

FIELD OF THE INVENTION

The field of invention relates to networking, generally, and, more
5 specifically, to link redundancy.

BACKGROUND

09745689-12200
10 Framers are commonly used in networking systems. Figure 1 illustrates an exemplary framer. A network line such as networking line 102 or networking line 103 in Figure 1 couples a pair of networking systems (e.g., switches, routers, multiplexers, gateways, etc.) so that the pair of networking systems may communicate with one another. Examples of networking lines include fiber optic or copper cable. Note that for simplicity, physical layer devices (e.g., lasers/photo-detectors, line drivers, etc.) typically placed between the framer 101 and the networking lines are not shown in Figure 1.

15 Frames (such as frames 108, 109, and 110) are used to organize the flow of information over a network line. In the case of SYNnchronous Optical Network (SONET) frames, each frame may be viewed as carrying "n" synchronous payloads envelopes (SPEs) of 810 bytes. Conceptually, as seen in Figure 1, frame 108 carries SPEs 111₁, 111₂, 111₃, . . . , 111_n; frame 109 carries SPEs 112₁, 112₂, 112₃, . . . , 112_n; and frame 110 carries SPEs 113₁, 113₂, 113₃, . . . , 113_n. For simplicity, note
20 that the overhead portions of each frame 108, 109, 110 are not shown.

For SONET frames, the time consumed by each frame (e.g., time T1 for frame 108, time T2 for frame 109, and time T3 for frame 110) corresponds to 125 μ s regardless of the number of SPEs carried per frame (i.e., "n").

Furthermore, the number of SPEs carried per frame remains constant from frame to frame. Thus, the number of SPEs carried per frame is indicative of the network line speed.

For example, a SONET networking line having only one SPE per frame (i.e., n=1) corresponds to a line speed of 51.840 Mbs (i.e., 810 bytes every 125 μ s).

Similarly, a SONET networking line having three SPEs per frame (i.e., n=3)

corresponds to a line speed of 155.52 Mbs (i.e., 2430 bytes every 125 μ s), a

SONET networking line having forty eight SPEs per frame (i.e., n=48)

corresponds to a line speed of 2.488 Gb/s (i.e., 38880 bytes every 125 μ s), etc.

Note that if the applicable networking line is optical "OC" is typically used instead of "STS" (e.g., OC-3, OC-48, etc.).

One SPE per 125 μ s is referred to as an STS-1 signal. Thus, a 51.840 Mbs SONET networking line carries a single STS-1 signal; a 155.52 Mbs SONET networking line carries three STS-1 signals; and a 2.488 Gb/s SONET networking line carries forty eight STS-1 signals. Typically, each STS-1 signal may be viewed as corresponding to the same SPE position across different frames. That is, a first STS-1 signal corresponds to SPEs 111₁, 112₁, and 113₁; a second STS-1 signal corresponds to SPEs 111₂, 112₂, 113₂; etc.,.

Figure 1 shows a framer within a networking system 110 that acts as a node in a network. The framer 101 in Figure 1 is one or more semiconductor chips that provide framing organization for a network line. For example, the exemplary framer 101 of Figure 1: 1) formats STS-1 signals into frames that are transmitted on an outbound networking line 103 to another network node; and 2) retrieves STS-1 signals from frames received from another network node on an inbound networking line 102.

In the case of outbound transmission, other portions of the framer's networking system 110 individually provide each STS-1 signal carried by the outbound network line 103 to the framer 101. For example, a first STS-1 signal is presented to the framer at input 107₁, a second STS-1 signal is presented to the framer at input 107₂, etc. Consequently, for example, the framer 101 maps on outbound networking line 103: the STS-1 signal received at input 107₁ across SPE positions 111₁, 112₁, 113₁; the STS-1 signal received at input 107₂ across SPE positions 111₂, 112₂, 113₂, etc.

Correspondingly, in the case of inbound transmission, each STS-1 signal carried by the inbound network line 102 is individually presented by the framer 101 to higher layers of the framer's networking node 110. For example, a first STS-1 signal mapped on SPE positions 111₁, 112₁, 113₁ is presented on framer output 106₁, a second STS-1 signal mapped on SPE positions 111₂, 112₂, 113₂ is presented on framer output 106₂, etc. The individual outbound STS-1 signals may be collectively referred to as outbound STS-1 signals 105. Similarly, the

individual inbound STS-1 signals may be collectively referred to as inbound STS-1 signals 104.

Note that different types of framers may exist. In one respect, the granularity of the inbound and outbound signals 104, 105 may vary. For example, with respect to SONET framers, each of the individual signals that collectively form the inbound and outbound signals 104, 105 may be comprised of a signal that consumes less bandwidth than an STS-1 signal (e.g., down to a 64kbps signal) or more bandwidth than an STS-1 signal (e.g., each input signal may correspond to a group of STS-1 signals such as an STS-3 rate signal or an STS-12 rate signal, or higher). In another respect, different framing formats may be applied (e.g., SDH).

002222 " 68954260

SUMMARY OF INVENTION

A method that changes the selection of a 2:1 multiplexer that receives a first output signal from a first framer and a second output signal from a second framer. The first output signal is the same as the second output signal.

002222 122200 09745689

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not limitation, in the Figures of the accompanying drawings in which:

Figure 1 shows an exemplary framer;

5 **Figure 2** shows an exemplary framer logic unit;

Figure 3 shows a 1+1 FFP PG or UPSR redundancy scheme that employs the exemplary framer logic unit of Figure 2;

Figure 4 shows an example of 1:N redundancy or 1:1 BLSR redundancy implemented with the framer logic unit 200 shown in Figure 2; and

10 **Figure 5** shows the exemplary framer logic unit of Figure 2 further including a switching or routing extension.

DETAILED DESCRIPTION

A method is described that changes the selection of a 2:1 multiplexer that receives a first output signal from a first framer and a second output signal from a second framer. The first output signal is the same as the second output signal.

5 An apparatus is described having a framer and a 2:1 multiplexer that receives an inbound signal from the framer. A first multiplexer receives at least one signal from another framer and the 2:1 multiplexer has an input coupled to an output from the first multiplexer. A second multiplexer receives at least one signal from the other framer and the second multiplexer has an output coupled to an input of the framer for an outbound signal.

10 ~~Link redundancy is a technique for protecting against the failure of a network line. A networking line may fail for any of a number of reasons (e.g., the line itself may be opened, the aforementioned physical layer devices may fail, etc.). As such, networking service providers and networking system providers are interested in technology that allows for such failures without disrupting the operation of a network.~~

15
20 Link redundancy is the notion that a "spare" network line may be installed into the network for the purpose of carrying another network line's traffic should the other network line fail. That is, if a network line fails, the network "switches over" to the spare network line in order to avoid significant disruption of the network.

Figure 2 shows an exemplary framer logic unit 200. The exemplary framer logic unit 200 of Figure 2 provides for efficient link redundancy, as described in more detail with respect to Figures 3 and 4. Before discussing how the architecture of the framer logic unit 200 of Figure 2 provides for efficient link redundancy implementation, however, the design of the framer logic unit 200 will first be discussed.

A framer 201 is shown coupled to an inbound networking line 202 and an outbound networking line 203. The framer 201 also provides inbound signals 204 from the inbound networking line 202 and receives outbound signals 205 for transmission over outbound networking line 203. In an embodiment, each signal corresponds to an STS-1 signal; however, other framer and signal embodiments may exist as alluded to in the background above.

Each inbound signal $206_1, 206_2, 206_3, \dots, 206_n$ is directed from the framer 201 to its own corresponding 2:1 multiplexer $208_1, 208_2, 208_3, \dots, 208_n$. That is, as seen in Figure 2, inbound signal 206_1 is directed to 2:1 multiplexer 208_1 ; inbound signal 206_2 is directed to 2:1 multiplexer 208_2 ; etc. The output of the framer logic unit 200 corresponds to the collective output 209 of each of the 2:1 multiplexers $208_1, 208_2, 208_3, \dots, 208_n$. Note that "n", as described in the background, in an embodiment, may correspond to each STS-1 signal carried per frame. Thus, for example, in an embodiment there is forty eight 2:1 multiplexers for a framer 201 that corresponds to an OC-48 framer (i.e., $n=48$).

09745689 12200

The framer logic unit output 209 may be viewed as a bus that collects each 2:1 multiplexer output. The framer logic unit 200 may be viewed as an architectural building block for a networking system. In such an application, the framer logic unit output 209 may be directed to every other framer logic unit in the system. That is, the architecture of the system may be designed to have multiple framer logic units (e.g., one framer logic unit 200 per inbound/outbound line pair coupled to the system).

As such, the framer logic unit 200 is designed to communicate with other framer logic units within the same system. For example, in an embodiment, framer logic unit 200 corresponds to a line interface card (LIC) that plugs into the backplane of a switch. In this case, the framer logic unit output 209 corresponds to a LIC output 209 that may be routed (through a backplane) to every other LIC in the system having the same (or similar) framer logic unit 200 design seen in Figure 2.

As such, the framer logic unit 200 is also designed to receive each output from the other framer logic units in the system. Thus, for a system having "x" other framer logic units, framer logic unit 200 receives "x" inputs $212_1, 212_2, 212_3, \dots, 212_x$. Thus, each input $212_1, 212_2, 212_3, \dots, 212_x$ corresponds to the output bus (similar to output 209) of another framer logic unit in the system.

Combining each input $212_1, 212_2, 212_3, \dots, 212_x$ together corresponds to a total input of "xn" signals. That is, in this particular embodiment, as each input has n signals and as there are x inputs, the total number of signals presented to

the framer logic unit 200 corresponds to "xn". The "xn" group of input signals are directed to a pair of (xn):n multiplexers 210, 211. Each (xn):n multiplexer 210, 211 selects "n" of the "xn" input signals.

Each of the 2:1 multiplexers $208_1, 208_2, 208_3, \dots, 208_n$ receives one of the n
5 outputs from the first (xn):n multiplexer 210. As described in more detail below with respect to Figures 3 and 4, the architectural arrangement of the first (xn):n multiplexer 210 feeding each of the 2:1 multiplexers $208_1, 208_2, 208_3, \dots, 208_n$ helps provide for efficient link redundancy.

The output of the second (xn):n multiplexer 211 corresponds to outbound
10 signals 205 that are transported over the outbound networking line 203. Note that the second multiplexer 211 may be used to establish the switching fabric of the networking system. That is, as the output signals of every framer logic unit may be presented to the second multiplexer 211, any input signal can be directed to the input of framer 201. As every framer logic unit may receive the output of
15 the other framer logic units within the system, an entire switching fabric for a networking system may be established by configuring the selection of particular output signals from each (xn):n multiplexer that feeds input signals to a framer.

Figure 3 shows a 1+1 Fiber Facility Protection (FFP) or Unidirectional Path Switched Ring (UPSR) redundancy scheme that employs the exemplary framer
20 logic unit of Figure 2. Under a 1+1 FFP redundancy scheme or a UPSR redundancy scheme, a "working" pair of networking lines (e.g., networking line pair 302a, 303a) are considered the primary networking lines used for

communication between the pair of nodes 370, 380 that are coupled by the pair.
Should the working pair of networking lines 302a, 303a fail, however, a
“protection” pair of networking lines 302b, 303b are enabled.

A networking system built with framer logic units 200 as discussed above
5 allows for efficient 1+1 FFP or UPSR redundancy because redundancy “hooks”
are built into the framer logic unit design. Specifically, note that Figure 3 shows
a first node 370 and a second node 380 that are each constructed with the framer
logic units 300a, 300b, 300c, and 300d as described above with respect to Figure 2.

Networking system 370 includes framer logic units 300c and 300d while
10 networking system 380 includes framer logic units 300a and 300b. Note that the
framer logic units within the same system are coupled as described above with
respect to Figure 2. That is, the output 309a of framer logic unit 300a is shown
coupled to an input 312b₂ of framer logic unit 300b. Similarly, the output 309c of
framer logic unit 300c is shown coupled to an input 312d₂ of framer logic unit
15 300d.

Note that the output 309b of framer logic unit 300b may be coupled to an
input of framer logic unit 300a and the output 309d of framer logic unit 300d may
be coupled to an input of framer logic unit 300c; however, for ease of drawing,
these couplings are not shown in Figure 3. Within the first networking system
20 370, the second (xn):n multiplexers 311c, 311d of each framer logic unit 300c,
300d are identically configured.

That is, both (xn):n multiplexers 311c, 311d present the same collection of signals to each of their respective framers 301c, 301d. As such, networking lines 302a, 302b are configured to send the same information from system 370 to system 380. Similarly, within the second networking system 380, the second

5 (xn):n multiplexers 311a, 311b of each framer logic unit 300a, 300b are identically configured such that the same information is sent from system 380 to system 370 over networking lines 303a and 303b.

002221" 6854760

10 The output of framer logic units 300a and 300c correspond to the inbound signals 304a, 304c from framers 301a and 301c, respectively. That is, the 2:1 multiplexers 308a₁ through 308a_x of framer logic unit 300a are configured to select the framer 301a inbound signals 304a rather than the output signals from the first multiplexer 310a. Similarly, the 2:1 multiplexers 308c₁ through 308c_x of framer logic unit 300c are configured to select the framer 301c inbound signals 304c rather than the output signals from the first multiplexer 310c.

Sub 72

15 The first (xn):n multiplexer 310b within framer logic unit 300b is configured to select the output signals from framer logic unit 300a. Also, the first (xn):n multiplexer 310d within framer logic unit 300d is configured to select the output signals from framer logic unit 300b. As such, first multiplexer 310b effectively presents framer 301a inbound signals 304a to the 2:1 multiplexers 350

20 within framer logic unit 300b; and first multiplexer 310d effectively presents framer 301c inbound signals 304c to the 2:1 multiplexers 360 within framer logic unit 300d.

5 The 2:1 multiplexers 350, 360 of frame logic units 300b, 300c also receive the inbound signals 304b, 304d from their respective framers 301b, 301d. As such, the 2:1 multiplexers 350 receive the inbound signals 304b from framer 301b at one channel input as well as receive the inbound signals 304a from framer 301a at the other channel input. Similarly, the 2:1 multiplexers 360 receive the inbound signals 304d from framer 301d at one channel input as well as receive the inbound signals 304c from framer 301c at the other channel input.

10 Because networking lines 302a and 302b send the same information as discussed above, the 2:1 multiplexers 350 within framer logic unit 300b effectively receives the same information at both channel inputs. Similarly, because networking lines 303a and 303b send the same information as discussed above, the 2:1 multiplexers 360 within framer logic unit 300d effectively receive the same information at both channel inputs.

15 As such, the 2:1 multiplexers 350, 360 may easily toggle from the working networking line pair 302a, 303a to the protection networking line pair 302b, 303b if the working networking line pair 302a, 303a should fail. Specifically, during normal "working" mode, the 2:1 multiplexers 350, 360 are respectively configured to enable the inbound signals 304a, 304c received from framers 301a, 301c. If a failure is detected on networking line pair 302a, 303a the selection
20 performed by the 2:1 multiplexers 350, 360 is toggled to produce the inbound signals 304b, 304d from the framers 301b, 301d at output 309b and 309d (rather the inbound signals 304a, 304c from framers 301a, 301c).

Thus, regardless as to which network line pair is actually used, the output 309b, 309d from framer logic units 300b, 300d are "looked to" by the networking system 380, 370 to as the source of information from the other networking system 370, 380. As such, either networking system 370, 380 does not experience significant disruption.

Figure 4 shows an example of 1:N redundancy or 1:1 Bidirectional Line Switched Ring (BLSR) redundancy implemented with the framer logic unit 200 shown in Figure 2. In either approach, a spare networking line pair 402_{spare} , 403_{spare} replaces any networking line pair that fails within a group of networking line pairs 402_1 through 402_N , 403_1 through 403_N . Note that the framer logic units $400a_1$ through $400a_N$, $400a_{\text{spare}}$ and framer logic units $400b_1$ through $400b_N$, $400b_{\text{spare}}$ each correspond to the framer logic unit design 200 of Figure 2.

In the approach of Figure 4, when a particular line pair within the protected group (402_1 through 402_N , 403_1 through 403_N) fails, the second (xn):n multiplexer 411a, 411b within each spare framer logic unit $400a$, $400b$ are configured to select the outbound signals 305a, 305b that were sent over the failed networking line pair. Thus, for example, if networking line pair 402_1 , 403_1 fails, the signals selected for transmission over networking line 403_1 are presented to framer $401a_{\text{spare}}$ for transmission over networking line 403_{spare} . Similarly, the signals selected for transmission over networking line 402_1 are presented to framer $401b_{\text{spare}}$ for transmission over networking line 402_{spare} . The

2:1 multiplexers of the spare framer logic units $400a_{\text{spare}}$, $400b_{\text{spare}}$ are configured to select the outbound signals from their respective framers $401a_{\text{spare}}$, $401b_{\text{spare}}$.

Figure 5 shows another framer logic unit embodiment 500 that may be viewed as the framer logic unit 200 of Figure 2 further including a switching or routing extension 501. The inputs $212_1, 212_2, 212_3, \dots, 212_x$ to the framer logic unit 500 are coupled to a third $(xn):n$ multiplexer 503 that may select any inbound signal for switching or routing. Routing or switching engine 501 provides packet based (e.g., Internet Protocol (IP) based) switching (e.g., label switching) or routing. The routing or switching engine 502 may be a logic circuit; or a processor that executes software consistent with the routing or switching protocol(s) to be employed; or a combination of logic and processor.

The routing or switching engine 502 assembles packets from the signals $505_1, 505_2, 505_3, \dots, 505_n$ selected by a third $(xn):n$ multiplexer 503. Based on the destination of a packet (e.g., as indicated in the packet's header), the routing or switching engine 502 determines an appropriate outbound signal $506_1, 506_2, 506_3, \dots, 506_n$ that the packet should be forwarded to. The packet is then disassembled and sent over the appropriate outbound signal. Note that a second framer logic unit output 502 may be viewed as a bus having each of the switching or routing engine output signals $506_1, 506_2, 506_3, \dots, 506_n$.

In an embodiment, framer output logic output 502 is sent to every other framer logic unit in the networking system. As such the networking system can be configured to provide packet based routing or switching between any of the n

inbound signals 505₁, 505₂, 505₃, . . . , 505_n selected by the third (xn):n multiplexer 503 and any of the "n" outbound signals 506₁, 506₂, 506₃, . . . , 506_n within the networking system.

Note also that embodiments of the present description may be implemented not only within a semiconductor chip but also within machine readable media. For example, the designs discussed above may be stored upon and/or embedded within machine readable media associated with a design tool used for designing semiconductor devices. Examples include a netlist formatted in the VHSIC Hardware Description Language (VHDL) language, Verilog language or SPICE language. Some netlist examples include: a behavioral level netlist, a register transfer level (RTL) netlist, a gate level netlist and a transistor level netlist. Machine readable media also include media having layout information such as a GDS-II file. Furthermore, netlist files or other machine readable media for semiconductor chip design may be used in a simulation environment to perform the methods of the teachings described above.

Thus, it is also to be understood that embodiments of this invention may be used as or to support a software program executed upon some form of processing core (such as the CPU of a computer) or otherwise implemented or realized upon or within a machine readable medium. A machine readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine readable medium includes read only memory (ROM); random access memory (RAM);

magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

In the foregoing specification, the invention has been described with
5 reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

10 For example, even though the above description has described a networking system where every framer logic unit 200 couples "n" signals to a framer 201 and every first 210, second 211, and third 503 multiplexers correspond to a (xn):n multiplexer, other embodiments may be realized that deviate from this architecture.

15 For example, framer logic units within the same networking system may couple different amounts of signals to their respective framer (e.g., some framer logic units may couple "n" signals to their respective framer while other framer logic units may couple more than (or less than) "n" signals to their respective framer. For example, a framer logic unit coupled to a pair of higher speed
20 networking lines may present more signals to its respective framer than a framer logic unit coupled to a pair of lower speed networking lines.

Furthermore a framer logic unit does not necessarily need to receive (at first 210, second 211, and third 503 multiplexers) every inbound signal in the system. As such, in light of the comments above, first 210, second 211, and third 503 multiplexers may exhibit a varied selection ratio from framer logic unit to .

5 framer logic unit. That is, some selection ratios within a networking system may be configured at $(xn):n$. However, selection ratios other than $(xn):n$ may exist within the same system. As such, the routing or switching engine 501 may switch or route packets from/to more than n signals or from/to less than n signals.

00222T 68954260